2000 Command and Control Research and Technology Symposium
Building the Readiness Data Warehouse

Ms. Sue Tysor, N662
Command-in-Chief, U.S. Atlantic Fleet
1562 Mitscher Avenue
Norfolk, VA
757.836.5523
757.836.9486
tysorsa@clf.navy.mil

Abstract

Regardless of the product or service being offered, a corporation, agency or department needs to thoroughly understand its customers and constituents. An organization needs to know how well it is executing its mission and how it can improve service. Organizations must manage costs as well as human and capital assets. The ability to fully leverage information assets can have a dramatic influence on each of these areas. To meet the requirement corporate data must be analyzed, comprehended, transformed and delivered. This is the role of the data warehouse. The data warehouse will deliver business intelligence based on operational data, decision support data and external data to all business units in the organization.

The program managers and resource managers at CINCLANFLT need consistent and reconciled business intelligence to manage the level of readiness of active and reserve forces. As a component command and force provider, readiness information is used by the FLTCINCs to make business management decisions about which assets should be used for assigned missions to support the national defense objectives of the country. As resources have gotten more scarce, the need for better readiness related information has increased. At the same time, many difficult questions have arisen which need to be answered, such as, how much readiness is the right amount and how do you measure it?

Our goal is to build a readiness data warehouse that will enable information management to change the way organizations leverage and value their information assets. With the ability to easily access information, mission delivery, resource management and data dissemination can be raised to levels previously unimagined.

This paper will identify the many issues associated with the process of building our readiness data warehouse. Specifically, it addresses the need to manage complexity and presents the development methodology, data architecture and technical architecture for the readiness data warehouse.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **JUN 2000**	2. REPORT TYPE	3. DATES COVERED **00-00-2000 to 00-00-2000**
4. TITLE AND SUBTITLE **Building the Readiness Data Warehouse**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **U.S. Atlantic Fleet,1562 Mitscher Avenue,Norfolk,VA,23551-2487**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**		**18**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

1. Introduction

This is one of the most challenging periods in recent times for DoD. With shrinking budgets, level or increasing tasking and a strong economy, which is successfully competing for personnel resources, it requires military decision makers to make more judicious use of all available resources.

Readiness of operating forces to carry out their assigned missions should be the goal of the entire service. Managing readiness to meet required goals is being achieved. However, it appears that it is being done at the expense of shore-based support. Both individual shore-based readiness statistics and the increased difficulty in making carriers ready for deployment support that conclusion. If true, this is not a situation that can be sustained for very long. To determine ground truth, business intelligence, which could be available with a data warehouse would be invaluable in determining the solution space for this problem and related resource management problems.

A specific example of the need for the readiness data warehouse comes from the AMSR study group, "the inability to control AVDLR costs within the context of the current budget limitations necessitates finding balance between readiness accounts and modernization/re-capitalization accounts. The difficulty experienced in accurately forecasting the FHP cost leads to under-funding/under-pricing, which then cause fleet under execution, bow-waving AVDLR into the next year and declining readiness." Having metrics and reliable detailed data, integrated in a data warehouse would allow for better management of costs by providing visibility of the consolidated picture of enterprise data.

The above states the compelling need for a data warehouse. The balance of this paper will focus on the building of the readiness data warehouse.

2. Objectives –What do we want to measure?

2.1 *Working with a Unique Mission*

The business of DoD is not exactly like that of any other company. There is not an example in the commercial world that has the same measures of performance, goals or missions across all business units as does the Navy or any other DoD service or agency. Therefore, we cannot take an existing model or product and apply it wholesale to meet our requirement. There are however strong conceptual similarities between commercial activities and DoD activities. For example, the functional areas of procurement, pay and benefits, inventory management, human resource management, information technology, and maintenance are areas where much can be learned from the successes and failures in the business world.   However, the fact remains, power projection and the ability to wage war and win a war are things only DoD does and how do you measure that.

2.2 *Changes in the Nature of Business*

After winning the cold war, the Department of Defense found itself beset with an ever widening array of missions to support. These were added to existing requirements being supported by the services. With the new missions, came new training requirements, new equipment, new environments and new rules of engagement, all of which may have new metrics and measures of performance to consider.

There is also the impact of new technology. This is being felt the most in the areas affected by changes in information technology. Changes in technology, generally bring changes in capability and with new or enhanced capability, comes a new wave of metrics to measure performance. For example,

The issue is that the readiness warehouse in its final state must be large enough handle all the relevant data needed to provide leadership with the business intelligence needed to make critical decisions. Another key aspect of the warehouse design is that it must be able to accommodate change because we cannot know all of the data requirements ahead of time.

2.3 *Managing the Complexity*

The task of building a data warehouse for an enterprise is a challenging undertaking. The complexity and size of a government service makes the task of building a readiness data warehouse is even more difficult. The difficulty of planning and implementing a single, undifferentiated, master data warehouse for the whole enterprise is monumental [Kimball et al., 1998]. The job is too overwhelming for most organizations and most mortal designers to contemplate. In support of the argument that data warehouses should be attacked incrementally, is the following excerpt from *The Data Warehouse Lifecycle Toolkit: Expert Methods for Designing, Developing, and Deploying Data Warehouses* which states, "The future of data warehousing is modular, cost effective, incrementally designed, distributed data marts. The data warehouse technology will be a rich mixture of large monolithic machines that grind through massive data sets with parallel processing, together with many separate small machines nibbling away on individual data sets that may be granular, mildly aggregated, or highly aggregated. The separate machines will be tied together with navigator software that will serve as switchboards for dispatching queries to the servers best able to respond."

2.4 *Top Down versus Bottom Up*

There are two competing philosophies regarding the basic approach to building the data warehouse. The "Top Down" approach requires that a completely centralized, tightly managed, single database be designed before any parts of it are summarized into individual Data Marts. Data Marts usually represent a subset of the overall data in the warehouse and built around a single business process or business unit. They are usually created for use by a specific functional department or customer group. The competing view is to build a warehouse from completely unrelated Data Marts.

Due to complexity of this task, time constraints and resources available it was decided that a hybrid approach would be used. The hybrid approach involves taking part of one philosophy and part of another to create a new approach. The new approach will hopefully retain the benefits of each parent and none of the detractors. The hybrid approach we have adopted accepts the fact that it is necessary to create an overall framework for the data warehouse. This is required to guide the design of each separate piece of the data warehouse. After you have designed such a framework, it is then possible to concentrate on the separate pieces, or data marts.

Another factor in the decision to use the hybrid approach is somewhat anecdotal but it is strongly supported by the authors, and that is, that the risk of failure is much higher when the scope of a project is very large. And the cost effectiveness of large projects is much lower than when a relatively small team, first develops an overarching framework and then concentrates on one, well-defined piece of the project at a time. The goal is to create a successful, repeatable process, which can be performed over and over until the project is completed.

2.5 *Targeting Mission Readiness*

Part of our plan to manage complexity is to focus on a subset of the overall data that will be needed for the readiness data warehouse. The questions that are asked most frequently at the higher echelon commands are related to the readiness of the war fighting entity. This is usually an individual activity, a Marine Expeditionary Force (MEF), Amphibious Readiness Group (ARG) or Carrier Battle Group (CVBG). These entities are all assigned missions. We want to measure each entities ability to perform the missions assigned. The right metrics will indicate the readiness level of each entity to perform one or more of it’s assigned missions.

Mission capability and readiness are also a large part of the requirements justification process. We must be able to review resource allocation based on mission and mission capability. This requires us to tie mission performance to mission cost.

Establishing a mix of metrics that clarify system goals, link decisions to goals and monitor processes for deviations will improve readiness. Connecting support element contributions to the overall system goals with appropriate ‘closed loop’ metrics will provide more rapid recognition of the causative relationship between support and readiness. Once cause factors are isolated through diagnostic measurements, resources can be focused to the correct areas. Focused support will more rapidly correct deviations from the goal and sustain readiness. The readiness data warehouse is a powerful tool to help solve these problems.

Readiness is a fundamental aspect of an armed force and can be viewed as the ability to rapidly mobilize, deploy and sustain trained forces in an area of operations to support specific missions, for an extended duration. Discussions of readiness components generally include the following six elements:

- Qualified people
- Combat-capable hardware and technology
- Appropriate levels of maintenance and spare parts for that hardware
- Appropriate tactics, techniques and procedures that support the capabilities represented by the qualified personnel and combat-capable hardware
- Training to ensure forces can actually conduct assigned operations
- The ability to deploy hardware and personnel to the fight

In order to assess how ready the military forces are, the following criteria can be applied and assessments made based on the results for each mission area:

- For each mission area, compare the required numbers of qualified personnel against the numbers actually on hand and available.
- For each mission, determine whether adequate supplies and spare parts are on hand.
- For each mission, determine and monitor the type and amount of training.
- Determine the ability of the sustaining base and infrastructure to support either major operations or smaller-scale contingencies for extended periods.
- Identify whether DoD has developed and promulgated the appropriate Operation Plan/Operation Order for conducting military operations.
- Determine the extent to which bases, hangars, maintenance depots, fuel farms, training ranges are in an “ready status”.

2.6 *CVBG Critical Tasks Concept Model*

The CVBG is assigned many critical missions or tasks.[1] These include Air Dominance, Power Projection, Maritime Superiority, Command and Control, Insert Land Forces, TBMD, Special Ops, Mine Warfare, Amphibious Ops, Combat SAR, Peacetime Presence, Sustainment and Surveillance and Intelligence. These missions are composed of critical tasks and each task is supported by sub tasks. The sub tasks are then linked to resource requirements, which are supported by data. The data is summarized to provide

[1] OPNAVISNT 3501.316

metrics. The metrics are then evaluated against performance goals associated with each mission. Figure 1 provides an overview of this process.

3. Development Methodology Process

As discussed earlier, we made a decision to breakup the task of building the readiness data warehouse into digestible parts. This could be done provided we first design the overall framework for our readiness data warehouse and then focus on the creation of the data mart.

3.1 *Data Warehouse Framework*

In developing the framework for the readiness data warehouse we had to examine the different options available for modeling data. The two data models that we evaluated were the entity/relationship model, which is most common to relational databases and a new discipline, which is referred to as the dimensional model. The dimensional model contains the same data as the entity/relationship model but the data is packaged in a symmetric format that makes the data more understandable by the user, enhances query performance and is easier to change.

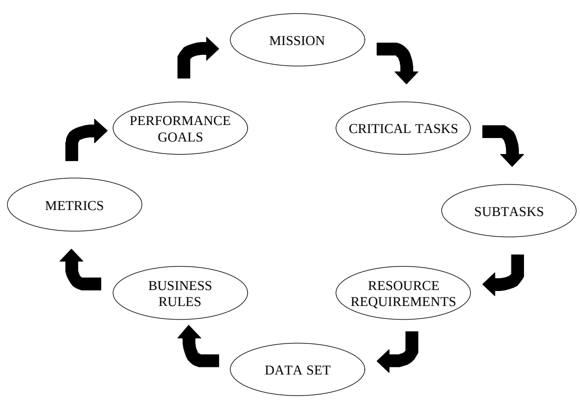

Figure 1. The CVBG Readiness Data Warehouse System Process Model

3.2 *Dimensional Model*

The traditional entity/relationship model seeks to normalize data. An unintended consequence of normalization is the loss of understandability for the customer and degraded performance. The dimensional model is an alternative to the entity/relationship model. It is designed to promote understanding and improve performance for the customer.

The main components of the dimensional model are fact tables and dimension tables. A fact table is the primary table in each dimensional model that is meant to contain measurements of the business. A fact, in this case, is a business measure such as the *number of tomahawk cruise missiles onboard*. Every fact table represents a many-to-many relationship and every fact table contains a set of two or more foreign keys that join to their respective dimension tables.

The dimension table is one of a set of companion tables that support a fact table. These are linked via a primary key, which supports referential integrity. Most dimension tables contain many textual attributes that are the basis for constraining and grouping within data warehouse queries.

3.3 *Single Physical Definition of an Attribute*

Different source systems that support the readiness data warehouse have evolved different lengths and data types for the same data element. It is essential, in building the readiness data warehouse that we use meaningful lengths and data types and these specifications must be consistent throughout the data warehouse. That is, all data marts within the data warehouse must be built from conformed dimensions and conformed facts.

3.4 *Consistent use of Entity Attribute Values*

All attributes in the data warehouse need to be consistent in the use of predefined values. Because many of the source systems use different attributes to represent the same meaning, these values need to be converted into a single, user friendly value as the data is loaded into the data warehouse.

3.5 *Issues Associated with Default and Missing Values*

A very real problem with building the readiness data warehouse is that the data being brought into the data warehouse is sometimes incomplete or contains values that cannot be transformed properly. This requires the transformation process to use well thought out, intelligent default values for missing or corrupt data. It is also important to provide visibility of defaulted data in the data warehouse. End users need to know the population of data they are using.

4. Readiness Data Warehouse Building blocks

The readiness data warehouse system is composed of several basic elements. Those basic elements are illustrated in figure 2.

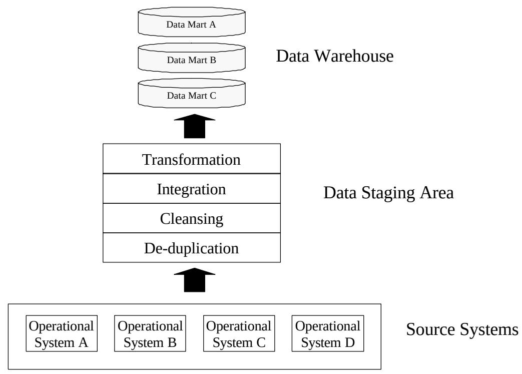

Figure 2. Basic Elements of the Readiness Data Warehouse System

4.1 *Separating Source Systems from the Data Warehouse*

One of the primary concepts of data warehousing is that data stored for business analysis can most effectively be accessed by separating it from the data in the operational systems [Gupta, 1997]. One reason for separating the data warehouse from the operational systems is that we needed minimize the impact on the operational systems. Another reason is that the data warehouse will bring in data from more than one operational system. This necessitated that the data be integrated at a place other than on the operational system.

The fundamental requirements of the operational system and the analysis system are different. The operational system is designed to capture the transactions of the business and therefore needs a very high mean time between failure and high availability. Our assumption is that source systems are not normally queried in broad and unexpected ways. Conversely, business analysis processes, supported by the data warehouse, are difficult to predefine and rarely need to have rigid response time requirements.

4.2 *Integrating Data from Multiple Sources*

The primary reason for combining data from multiple source systems is the ability to cross-reference the data. Nearly all data in a typical data warehouse is built around the time dimension. Time is one of the primary filtering criterion for analysis within the data warehouse. For example, an analyst may want to generate queries for a specific week, month, quarter or year. Or one may compare year-on-year activity.

The readiness data warehouse serves as an effective platform to merge data from multiple business applications. It can also integrate multiple versions of the same application and

it can allow for year-on-year analysis even though the base operational application has changed [Gupta, 1997].

4.3 *Source Systems*

Source systems are often referred to as operational systems or "legacy systems". For the readiness data warehouse, source systems are everything from batch loaded, IBM mainframe systems which house corporate personnel data to event by event updates to regional readiness databases holding maintenance and supply data to on-line-transaction-processing, OLTP, systems which capture loosely aggregated mission capability data such as fuel, weapons and food stores.

The wide variety of source systems and multitude of communication processes found in our environment further complicates the task of building our readiness data warehouse. Added to this are the data quality issues associated with the data stored in most Navy readiness databases. Without reliable, accurate data, it is not possible to make the best management decisions regarding such thing as fuel allocations, maintenance decisions and deployment decisions. The dollars represented by a decision on fuel can represent many millions of dollars. The problem can be summed up by Heilman's axiom, "You can't manage what you can't measure; and you can't measure what you can't define".

The data problems are caused by many factors including poor data fidelity, inaccurate data, improper reporting formats, multiple conflicting sources and difficult reporting processes. As a result, many critical management or tactical decisions about the employment of scarce Fleet resources could be negatively impacted. The issue of redundant metrics and conflicting data sources is exemplified by issue number 3 from the Aviation Maintenance and Supply Readiness (AMSR) working group, "Several different ILS data collection improvement efforts are currently under development by the Air TYCOMs and NAVAIR 3.0 as interim tools while awaiting NALCOMIS optimized to be fielded starting in mid FY99. These efforts should be refocused and integrated into a single interim ILS metric initiative able to provide leadership at all levels with real-time, holistic, end-to-end insight into an operating unit's logistic health." Again, from the AMSR working group, issue number 4 addressed Data Integrity Improvement, the system "fails to provide total data visibility to resource manage. During FY-97, COMNAVAIRPAC lost 30 to 40 percent of aviation 3M maintenance data submitted. The loss of this data has a substantial impact on the information being gleamed from aviation 3M data for resourced management decision making. Corrective action is needed to accurately capture and properly record all data submissions.

Accurate, consistent and timely readiness reporting is essential to monitoring the level of readiness of active and reserve forces. As part of the data warehouse effort, a hub-and-spoke framework or infrastructure approach is being used to support the data warehouse project. This is contrasted with the more traditional approach of using non-integrated, point solutions. While point solutions can often allow you to implement individual pieces of the data warehouse architecture successfully, they tend not to be integrated which makes maintenance a costly nightmare and responsiveness impossible.

There are two types of entities at the end of the spokes: source systems that feed source data to the warehouse and user-oriented systems that are fed data from the warehouse.

One major objective is to avoid imposing large data extraction programs on the data sources. Instead, programs are built to push the data to the hub. The hub then handles the transformation, loading and archival of the data. In addition, the hub manages the dispatching of data to end-user data marts, OLAP tools and directly to user tools. This allows for the monitoring and maintenance of data through the entire information supply chain.

This system design is being process driven and the adherence to processes will impose rigor on the entire system. Keys to a successful well-defined process are that the process must be repeatable, have a process owner and captures data as close to the source as possible.

Another objective is to fix the data at the source and not to repeatedly fix the data at the back end. Lastly, it is imperative to ensure that the cleansed data satisfies the users and results in the ability of the users to make mission-critical decisions.

It is also essential that there be one authoritative source identified for all data and that data flow from the originator to the user quickly and accurately. This model eliminates database inconsistencies (different values in different databases for the same element), automates data input which helps to eliminate database data from being out of date because people have not manually entered new data (a common failing of databases) and breaks down "stovepipe" systems implementing system feedback loops.

4.4 *Data Staging Area*

The data staging area refers to everything between the source systems and the data warehouse. This is where the data is organized, cleansed, integrated, transformed and archived for use in the data warehouse. In our case, the data is manipulated in a normalized structure in a relational database. However, it will not be used to provide support for queries or presentation services.

The issues associated with the logical transformation of data brought from the source systems to the readiness data warehouse required extensive planning and design effort. This is probably one of the most critical steps in the development process. The design must be efficient and flexible to be able to accommodate all of the business data from many different source systems. The term flexible refers to the ability of the system to be extensible so that data from new applications can be added when a business case can be made for adding the data.

Our readiness data warehouse model aligns with the business structure of the Navy. For example, the Navy is organized around the war fighting structures such as, CVBG, MEF and ARG and warfare communities such as submarines, airplanes and surface ships. The

Navy also manages by resource area such as maintenance, supply, personnel and training. Figure 3 illustrates the alignment of data warehouse entities with the business structure.

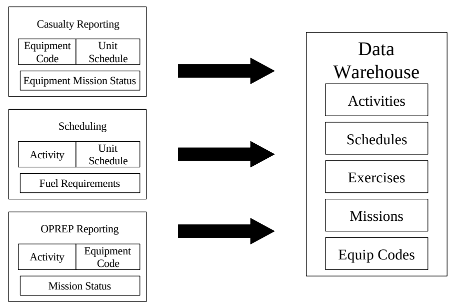

- No data model restrictions on source systems
- Data warehouse model has business entities

Figure 3. Readiness Data Warehouse Entities Align with the Business Structure

4.5 *Data Transformation*

Data transformation is necessary to clean up the data coming in from the source systems.

The conceptual, physical data transformation process is illustrated in figure 4.

Another major role of data transformation from source system to data warehouse is all about making the data useful. For this reason, many of the terms used in the source systems are transformed into standard business terms. The successful readiness data warehouse will use standard business terms that are self-explanatory.

4.6 *Operational Data Store*

The operational data store was originally defined as a frequently updated, volatile, integrated copy of data from operational systems that is meant to be accessed by "clerks

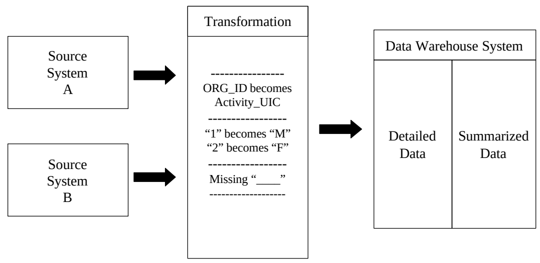

Figure 4. Conceptual Physical Data Transformation Process

This concept has been replaced by a concept that incorporates the operational data as a full participant in the data warehouse with performance-enhancing aggregations and associated time histories. It has been referred to as the “front edge” of the data warehouse.

This is important in our view because we do store large quantities of detail data in our readiness data warehouse. We take many snapshots, for example, of the changing statuses of supply parts that have been requisitioned for surface ships and tactical fighters. Tracking the statuses over time provides very useful information about business processes that support those efforts.

4.7 *Presentation Server*

This is the physical machine, which holds the data warehouse data. This data is organized to support direct querying by the end users to produce reports, graphs and other applications. We have control over the data model in the presentation server and it is here that the data is stored and presented in a dimensional framework to support the end user. This is the analyst’s source of data for the enterprise.

A key aspect of the presentation server is that it performs summarization and pre-defined analysis of data. These summary business views are often generated by summarizing detail data and applying business rules to the detail data. These rules can be very complex and they can support many different views into the same data. The summary views can hide all the complexities from the end user making the system more user friendly. Only analysts that perform data mining need to understand warehouse detail records and all the business rules.

4.8 *Data Mart*

The data mart is a subset of the readiness data warehouse. It usually reflects one business process or supports a homogenous group of business customers. It conforms to the overall data warehouse dimension framework and can be developed it its entirety without hindering subsequent data mart development. Finally, data marts contain granular data and may or may not contain performance enhancing summaries [Kimball et al., 1998].

4.9 *Data Warehouse*

The readiness data warehouse is the source of readiness data for the enterprise. It serves as focal point for the union of all the data marts.

One of the primary goals of the data warehouse is to make it as flexible and accessible as possible. For this purpose, there are many tools available to support use of the warehouse. These include things from simple query engines to multi dimensional analysis tools. There is an interesting lesson we have learned from working with relational databases and the readiness data warehouse, and that is that most users tend to want to get the same information out of the new warehouse that they were able to get using the old tools. There is apprehension about using the new system for more than generating the same reports they always did. It is only after they start to have significant input into the development process that they become advocates of the new capability and champion the new system. It is for this reason that most tools that get used initially are on the low end.

It is also a feature of the readiness data warehouse that it does not contain operational state information. Data in the source systems can be very dynamic and constantly going through state changes. Although the state changes may be recorded in the data warehouse the dynamic nature of the data does not. The result is that we will carry many periodic snapshots of operational states of certain data into the data warehouse. For this reason, loading the data warehouse is controlled as data is corrected and statuses and labels are changed.

The Readiness Data Warehouse consists of data schemas (arrangement of tables and table joins) for readiness data. Within each of these schemas are de-normalized tables, accompanied by a star or snowflake schema of a normalized fact table joined by de-normalized dimensional tables.

The Readiness Data Warehouse interface requires a quick response to a user's request for data. Therefore, to eliminate the delay of multiple table joins needed to create the record set, flat de-normalized tables are created which provide little processing of data. Using industrial terminology these tables are the Data Marts. De-normalized tables contain all of the data with as few joins a possible. An example of a de-normalized table is shown in Figure 5. An example of a normalized schema is shown in Figure 6. As seen in this figure, the main table is normalized while the look-up tables are de-normalized.

An example of the high level system architecture for the readiness data warehouse system is illustrated in Figure 7.

4.10 *On-Line Analytical Processing*

The readiness data warehouse will also support On-Line Analytical Processing (OLAP). The On-Line Analytical Processor (OLAP) utilizes star or snowflake schemas to conduct its processing. Normalized tables, are used to reduce the storage space of the table. Therefore,

UIC	Tycom	HomePort
N30720	CNSL	YORKTOWN
N03362	CNAP	YOKOSUKA
N55178	CNARF	WILLOW GROVE
N08981	CNARF	WILLOW GROVE
N09172	CNAL	WILLOW GROVE
N09174	CNAL	WILLOW GROVE
N09616	CNAP	WHIDBEY ISLAND
N09632	CNAP	WHIDBEY ISLAND
N09674	CNAP	WHIDBEY ISLAND
N09707	CNAP	WHIDBEY ISLAND

Figure 5. De-Normalized Table

where applicable, the information in a column contains an identification number (ID) and points to a look-up table containing the appropriate data. These Ids are typically integers that take less table space to maintain.

5. Data Warehouse Architecture

The architecture of the data warehouse and the data warehouse model greatly affect the success of the data warehouse. Figure 8 lists the pros and cons for adopting an N-Tier architecture.

5.1 *N-Tier Architecture*

The system architecture supporting the readiness data warehouse is a three-tiered application based on Microsoft Windows Distributed InterNet Architecture (DNA). Refer to Diagram 9

UIC_ID	UIC
1	N30720
2	N03362
3	N55178
4	N08981
5	N09172
6	N09174
7	N09616
8	N09632
9	N09674
10	N09707

Tycom_ID	Tycom
1	CNSL
2	CNAP
3	CNARF
4	CNAL

HomePort_ID	HomePort
1	YORKTOWN
2	YOKOSUKA
3	WILLOW GROVE
4	WHIDBEY ISLAND

UIC_ID	Tycom_ID	HomePort_ID
1	1	1
2	2	2
3	3	3
4	3	3
5	4	3
6	4	3
7	2	4
8	2	4
9	2	4
10	2	4

Figure 6. Normalized Schema

below for an overview. The system architecture is segmented into three logical tiers of functionality that include: presentation services, business services, and data services.

5.1.1 *Business Services*

The responsibilities for testing the business logic tier includes receiving and checking the input from the presentation tier, interacting with the data services to perform the business operations that the application was designed to automate (for example, unit readiness, specific CASREPS, SORTS, and so on), and sending the processed results to the presentation tier for comparison. The tester will be testing specific logical functionality and queries.

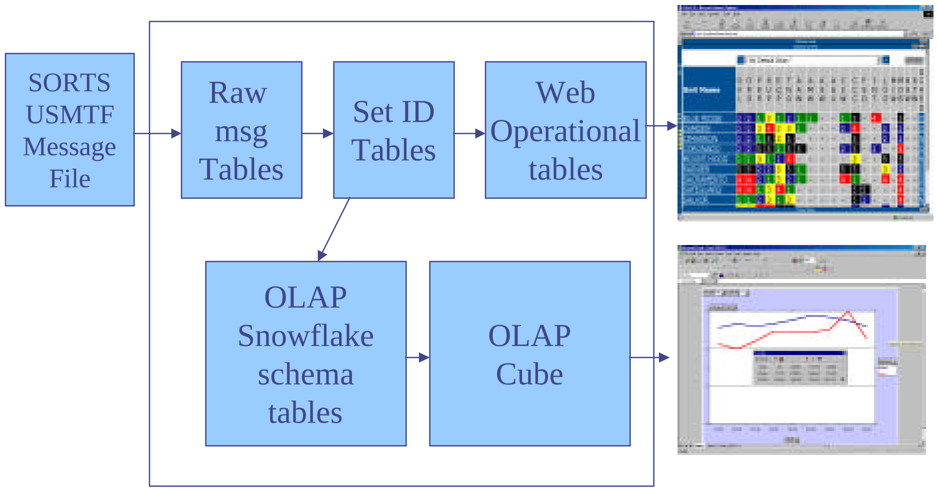

Figure 7. High Level System Architecture for the Readiness Data Warehouse

Pros	**Cons**
• Multi-Language Support	• Increased Network Traffic
• Centralized Components	• More Physical Resources
• Load Balancing	• Increased Complexity
• Efficient Data Access, including external resources	
• Improved Security	
• Scalability	
• Reliability	
• Mission Critical	

Figure 8. Pros and Cons for an N-Tier Architecture to Support the Readiness Data Warehouse

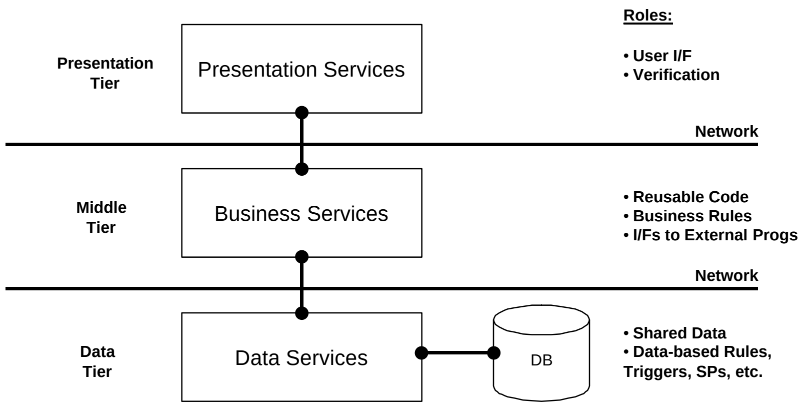

Figure 9. Distributed InterNet Architecture

5.1.2 *Data Services*

The responsibilities for testing the data tier include validating the storage of data, retrieval of data, maintenance of data, and integrity of data. The validation of the data will be tested for major test cases. Integrity of the data will be tested during times of disconnect and stop/start testing.

5.2 *The Role of Middleware*

Middleware is a concept that allows distributed, clients and databases, that are assembled from many different components, to interoperate seamlessly. This capability is often referred to as virtual application computing. This concept is important because many corporate applications running today, are application stovepipes. Stovepipe applications generally support a single domain, such as, maintenance, logistics, training, etc., but they usually evolve in isolation from one another and do not work well with applications from other domains.

There are two reasons for using middleware in support of the readiness data warehouse. The first reason is that new source systems are continually introduced and some of them contain data that support a business case for incorporating data from them. The second reason is that the CLF readiness data warehouse needs to interoperate with other data warehouses. Both of these reasons support the use of middleware to reduce the time and cost of building unique integration programs each time another system comes along.

We are currently incorporating support for XML in the belief that it offers a solid data exchange standard that improves interoperability between heterogenous data systems. XML is emerging as a strong contender to become the data interchange standard of the web. Data content is separated from its presentation format, allowing customized views of data tailored to support specific user requirements.

Figure 10, provides a conceptual example of an integrated readiness data warehouse environment.

6. Summary

The goal of building the readiness data warehouse is to develop a system that is easy to use, provides reliable, timely readiness information and is tightly integrated with all other data that is related to the readiness assessment processes.

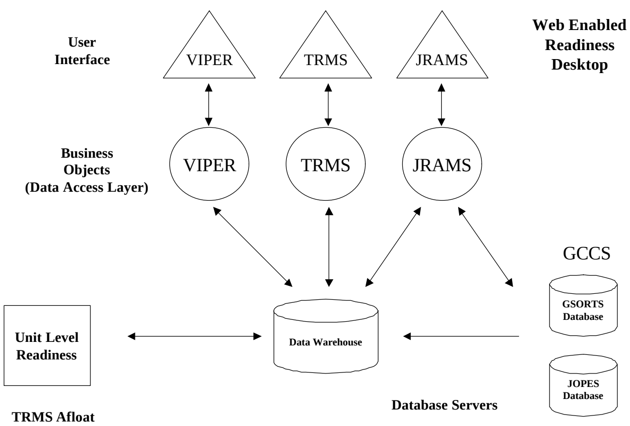

Figure 10. Conceptual Integrated Readiness Data Warehouse Architecture

The ultimate goal is to end up with a Force Planning capability that associates resource requirements with Force Plans that are developed to meet National Security objectives such as “presence” and to be able to do “what if” planning. Navy leadership desires to be able to quantifiably demonstrate requirements, impact and cost associated with Force level decisions. This tool would be able to support analysis necessary to prepare for and respond to events such as the Quadrennial Review.